



BANK OF ENGLAND

# Speech

---

## What's going on? Uncertain data and uncertain outcomes

Speech given by

Martin Weale, External Member of the Monetary Policy Committee

University of Liverpool

13 May 2016

I am grateful to Stuart Berry, Philip Bunn, Alan Castle, Shiv Chowla, Jonathan Fullwood, Marco Garofalo, Tomas Key, Jeanne Le Roux, Alan Mankikar, James Mitchell, Andre Moreira, Lee Robinson, Matt Swannell, James Tasker, Ken Wallis and Abigail Whiting for helpful discussions and support, and to my colleagues on the Monetary Policy Committee and at the Bank who provided valuable comments on an earlier draft.

*"I know what's going on. ... I'm going on"*. Rt. Hon. Harold Wilson MP, 4<sup>th</sup> May 1969.

## Introduction

Good afternoon and thank you for inviting me here to speak to you.

I thought that quoting Mr Wilson, a local MP as well as Prime Minister, in his centenary year would be an appropriate way of introducing a discussion of some of the practical difficulties faced by monetary policy makers. First, however, let me assure you that, unlike Mr Wilson, I am not going on; I have been lucky to have enjoyed three opportunities to visit Liverpool during my time as a member of the Monetary Policy Committee, but this will be the last. In early August I will be replaced on the Committee by Michael Saunders, from Citigroup, to whom I offer my congratulations.

Descriptions of the monetary policy making process often invoke various vehicle-driving analogies, with monetary policy decisions being seen as applications of the accelerator or brake that aim to keep the economy on track. Unfortunately, as many people have previously noted, these sorts of analogies are misleading for several reasons. The one on which I would like to focus today is that, compared to the drivers of most vehicles, the MPC is much less certain about how fast the economy is currently moving, or even how fast it was going in the recent past.

Put another way, unlike Mr Wilson, while we do our best to form judgements about what is going on in the economy, I would never make the claim that we "know" what is going on. There are always degrees of uncertainty about both the data and their interpretation. I would like to discuss today some of the means by which we try to minimise that uncertainty and also the ways in which we try to express the limits to our knowledge.

I will begin with some of the perennial challenges that the MPC faces before focussing on present sources of uncertainty. So first I will talk about how we try to estimate the state of the economy ahead of any official data. Secondly I will take a fresh look at the relationship between the first official estimates of GDP growth and final estimates of growth. I will follow this by offering my own thoughts on how the Committee might represent the uncertainty surrounding its forecast before proceeding to discuss some particular sources of uncertainty at present.

## What is the current state of the economy?

How fast the economy is currently growing might seem a simple question to ask. The answer is anything but straightforward. The Committee needs to make decisions as things happen, but data become available only with a lag. The Office for National Statistics is faster than most statistical offices at producing estimates of

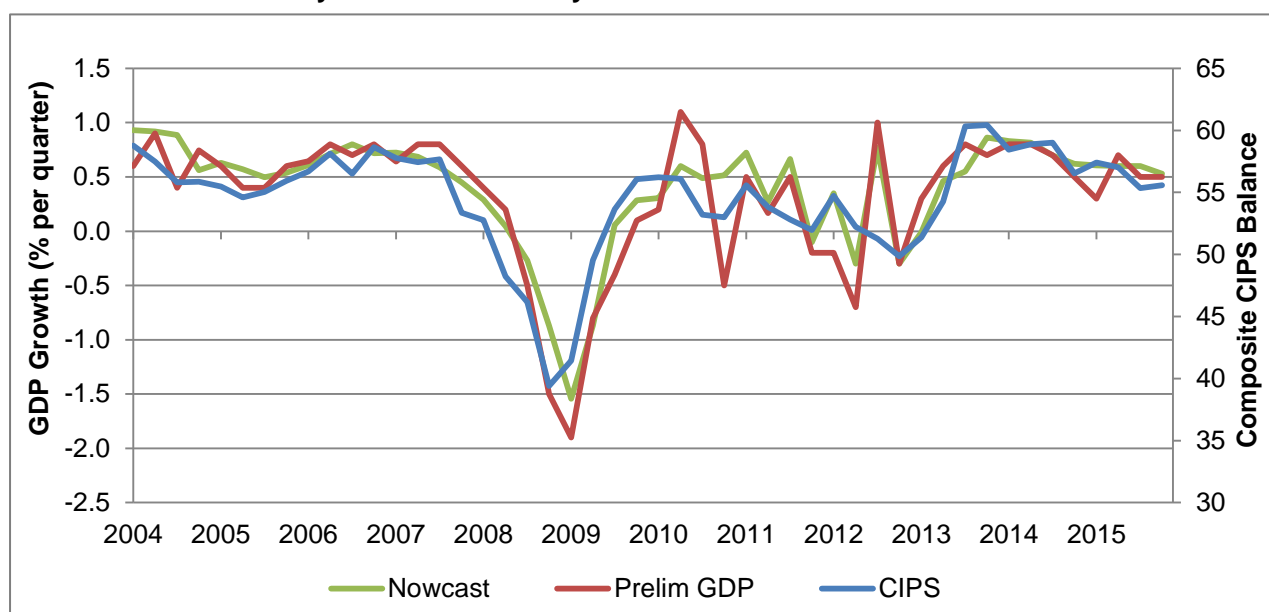
economic growth, but their first estimates appear only about four weeks after the end of the quarter to which they relate. And those estimates, like those produced by other countries, are subject to substantial revision.

In the short term, therefore, we face the problem of estimating the current state of the economy. 'Nowcasting', as this process is commonly known, the growth rate of GDP can be done both using forecasting methods – e.g. producing estimates of GDP on the basis of the most recent solid economic data available – and by using those data which become available during the quarter in question. Not surprisingly there is a preference for relying on current information as providing an indicator of what is going on at the moment and this comes in two forms: business surveys and the accrual of official monthly data.

Business surveys are a popular and widely-used type of current information on which a great deal of attention is focused. These are compiled by bodies such as the British Chambers of Commerce (BCC), the Chartered Institute of Procurement and Supply (CIPS) and the Confederation of British Industry (CBI). Typically a relatively small number of businesses are approached and respondents are asked to provide qualitative answers to a range of questions. The questions of most interest to the Committee are about the recent, and expected future, growth rate of firms' output. Attention usually focuses on a balance – the proportion of respondents who report an increase in output less those who report a decrease.

In fact surveys on their own do not do very well as current indicators (Mitchell, 2009) and the Bank uses a mix of models and judgement to produce its nowcasts (Bell et al., 2014). These make use of both survey information and monthly data as they accrue together with other information. In Chart 1 I show the Bank's nowcast together with the quarterly average value for the CIPS survey and the ONS preliminary estimate. You can see that the Bank's nowcast is, overall, a better estimate than could be inferred from the mean of the CIPS survey.

**Chart 1: The CIPS Survey and the Preliminary ONS and Bank Nowcast Estimates of GDP Growth**



Sources: ONS, Markit/CIPS and Bank calculations

The outperformance of the Bank's nowcast is, in fact, stronger than simple visual inspection suggests. The nowcast is published in each *Inflation Report*, which is produced less than half way through the quarter, while the survey measure, in contrast, is the average of the three monthly surveys and is not available until the end of each quarter. I currently take modest reassurance from the fact that our nowcast is for growth of 0.3 per cent in the current quarter, whereas the most recent CIPS figures have been widely quoted in the press as implying growth of 0.1 per cent. The nowcast estimate is, however, an estimate of the final figure and, as I will show later, outside the abnormal period of the recession of 2008/9, low initial estimates are prone to upward revision. So the initial estimate may well be lower.

On top of this, of course, over the period from 2004Q1 to 2015Q4 the root mean square error of the nowcast is 0.28 per cent, suggesting that there is only about a chance of two out of three that the outcome is within 0.28 per cent of the nowcast value. The margin of uncertainty is quite wide. I will not fall into the linguistic trap of suggesting that nowcasting is difficult – it is not. But it is inexact. We cannot be precisely sure of what is going on as we make our decisions and produce our forecasts. Indeed there has to be quite a high chance that the figure deduced from the CIPS data is in fact closer to the ONS preliminary estimate than is our own figure. There is equally, of course, quite a high chance that things are materially better than these figures suggest.

## **Could Business Surveys do Better?**

There is nevertheless a question of why business surveys do not do better as indicators of the state of the economy, or equally a question of whether they could be designed so as to do better. In order to explore this question my colleague, Tomas Key, and I have constructed a 'synthetic' business survey from the monthly data that firms provide to the Office for National Statistics, which it uses to compile its estimates of GDP.

We categorised each firm's output as up, no change or down on the month, and then compared the balance of these 'responses' – the proportion of firms reporting up relative to the proportion of firms reporting down – with the movements in output directly implied by the figures. This allows us to see how much information is lost by moving from a quantitative to a qualitative measure of output growth. The answer is that, with an optimally designed survey, the loss of information is minimal – the survey balance has a correlation with the actual data of only just less than one.

By an optimally designed survey, I mean a survey that asks firms whether their output has grown above a certain threshold, declined below a certain threshold or has grown by an amount in between these two. The thresholds that we have found to be optimal are quite wide, +/- 15%, and considerably wider than the thresholds of +/- 2-4% that firms themselves have typically told the CBI answering practices survey that they

are using.<sup>1</sup> If we examine the information loss from a survey with these thresholds, the loss is greater, but the correlation with the actual data is still around 0.9.

It therefore seems likely that it is not the fact that business surveys are qualitative that reduces the quality of the information that they provide, but other factors. One is that the only way that business surveys can be timelier than the ONS estimates is if they have smaller sample sizes, which will necessarily reduce the quality of the signal that they provide. Another is that firms often do not answer the surveys in the exact way that they are asked to. For example, Lui, Mitchell and Weale (2011) have found that firms' answers to the CBI's Industrial Trends Survey about growth in output over the last three months contain information from up to six months in the past.

There is, however, an additional factor which explains why surveys do not do better. They do not cover the output of the public sector, which is an important part of GDP and there may be some other sectors where coverage is relatively limited. So, while it may be possible to improve on their design, the gain in performance which it seems could be delivered by improving the surveys would take us only so far.

## **Preliminary GDP and Final GDP**

Producing an estimate of economic growth ahead of the first official estimate is only one stage in our process of trying to form a view as to what is going on. Official estimates of the national accounts, and in particular real GDP, are subject to considerable revision or, as I prefer to call it, updating.<sup>2</sup> This should hardly be a surprise. After all policy-makers like the members of the Monetary Policy Committee want data to be up to date. But data are collected from a range of sources and some of these become available only with a considerable delay.

As you know, GDP can be estimated from the output of businesses, the expenditure of households, firms and the government or from measures of income. These three different measures give different answers, and current thinking is that the most accurate estimate of GDP is obtained from measures of income computed from tax data. In the nature of things, however, these tax data are available fully only after people have completed their tax returns and this is often not done until the January following the end of the financial year in question. A first estimate of the output measure is available, in contrast, less than a month after the end of each quarter in question; even this is, however, subject to revision because the preliminary estimate is computed, as it needs to be, from an incomplete sample of returns. As later data are included the number is quite rightly revised. There are further revisions as adjustments are made to align the output and expenditure measures with income estimates as they become available.

---

<sup>1</sup> See Bush (2008). One reason for this difference is likely to be that the CBI, in its Industrial Trends Survey, asks respondents about their firm's output growth after subtracting seasonal factors, whereas we follow the CIPS methodology of 'asking' about growth including for seasonal reasons.


<sup>2</sup> A suggestion made by Sir Adrian Smith.

For the user of data, this then raises two questions. First of all, how good is the preliminary measure as a predictor of the final measure, taken here to be the estimate published after three annual cycles of revision? Secondly, how far can one predict the revisions to the preliminary estimate, and in particular, are business surveys of any help in that regard?

Chart 2 shows the preliminary estimate of growth plotted against the final estimate of growth for calendar quarters from 1997Q1 to 2012Q4, the last quarter for which we have a final estimate. I have also estimated the regression line linking the two, and this is shown in the chart. If the preliminary estimate is a good estimate of the final estimate we should expect the coefficient in the regression equation to be close to one and the constant term close to zero. That does not say that the numbers are perfect; revisions still happen. But it does suggest that there is no systematic bias in the initial data. The regression equation seems satisfactory, and a formal statistical test suggests that I cannot reject the hypothesis that the slope coefficient is one with a constant of zero – the p-value is 0.34.


Unfortunately that is not the end of the matter. The diagram shows a classic pitfall of careless applied work. I have plotted in red the observations for the period 2008Q3 to 2009Q2, the period of sharp economic contraction. You can see that three of these stand out very sharply from the rest of the pack. In such circumstances regression analysis can be highly misleading. The outliers exert a very great deal of influence on the estimates of the slope and constant; to be confident in the conclusions one would want to be sure that they did not depend on the inclusion of these recession data. You can see, in Chart 3, that when I leave out the four outlying quarters of the recession, a very different picture emerges. The slope coefficient is well below one and the constant term is materially positive. If I test the hypothesis that the slope is one and the constant is zero then, with these observations, the probability of that being true is less than 0.01 per cent; in other words, the statistics are saying that we should rule it out as being rather unlikely.

**Chart 2: GDP Revisions**



Sources: ONS and Bank calculations

**Chart 3: GDP Revisions Excluding Recession**



Sources: ONS and Bank calculations

I am, in effect, arguing that the measurement issues which arise during sharp contractions, like the period from 2008Q3 to 2009Q2 are different from those which arise in more normal times.<sup>3</sup> That might raise the question whether the new, post-recession revision pattern is different from that found before the recession. The answer to this seems to be no, or at least not materially. If I look at the period before the crisis, I find a slope of 0.70 and a constant of 0.20, while for the period after the crisis, the slope is 0.62 and the constant 0.25. A formal statistical test suggests that I cannot reject the hypothesis that these are equal – the p-value is 0.8. So, while many things are different now from before the crisis, it is not obvious that the relationship between preliminary GDP and final GDP has changed.

The stability of this relationship might come as something of a surprise. However, as I show in Appendix A, this is a natural consequence of producing estimates from incomplete data. They are likely to exaggerate movements in both directions. The framework that I set out in the Appendix assumes that the economy can be split into two sectors, sector A and sector B. The first estimates are produced using only data for sector A because those for sector B are not yet in. Plainly a better first estimate could be produced if, in addition to the hard data for sector A, a good forecast is used for sector B. On the other hand, if the forecast for sector B is not very reliable, then it may be better to rely just on the forecast for sector A and make an adjustment on the basis of the regression equation that I have set out above.

The approach used by the MPC to adjust ONS initial estimates (Cunningham *et al.*, 2012) makes heavy use of business surveys. I suppose I like to think of it as constructing a substantial infrastructure to forecast the output of sector B, at least when applied to the ONS preliminary estimate of GDP growth.<sup>4</sup> In practice of course, because we cannot see, in the ONS release, separate growth rates for sector A and sector B, it simply produces an estimate based on the initial estimate of GDP, together with past estimates and additional data provided by business surveys.

For today's purposes I would like to keep it simple. In Table 1, I show, as model 1, a regression equation estimated for the period before the crisis, with final GDP growth as the dependent variable and the preliminary estimate as the explanatory variable. Model 2 includes a composite balance constructed from the CIPS surveys as an additional explanatory variable. Neither t-statistic is significant at five per cent, pointing to a problem of collinearity; the survey data and the preliminary estimate are highly correlated. Model 3 is the sort of model which would be used by someone who believed that ONS first estimates were no good. It relies only on the CIPS balance and has a slightly larger standard error than model 1.

How well do each of these work in the period after the crisis? Model 1 gives an out of sample root mean square error of 0.20. Model 2 gives an error of 0.22 while model 3 gives 0.40 for the same period, which is

---

<sup>3</sup> There is also a case for removing the effects of revisions associated with definitional changes to GDP – I have not done so here.

<sup>4</sup> The Bank's model explores all vintages of GDP growth from the preliminary estimate to the final estimate, while here I am considering only the move from the preliminary estimate to the final estimate.

larger than the root mean square revision of 0.34. The root mean square error of the backcast published in the *Inflation Report* is 0.24, so, at least over this period, models 1 and 2 perform better.

That does not, of course, mean that the MPC's approach will always be worse than model 1 or model 2, or indeed that model 1 will always perform better than model 2. The conclusion I would draw from this is simply that I should view cautiously any claim that business surveys provide helpful information about the current state of the economy, over and above what the ONS can tell us in the preliminary estimate. At the same time, of course, it is even more important to remember that equations such as model 1 and model 2 do not provide the exact answer. The root mean square errors suggest that only about two thirds of the final estimates will lie within 0.2 percentage points of the backcasts that they can generate. We cannot be sure of what is going on, even when we have initial ONS estimates of economic growth.

**Table 1: Modelling Final GDP Growth**

Dependent variable is Final Estimate of GDP Growth	Model 1	Model 2	Model 3
Preliminary GDP Growth	0.70 (5.14)	0.39 (1.79)	
CIPS Composite Balance		0.03 (1.74)	0.06 (5.11)
Constant	0.20 (2.31)	-1.33 (1.50)	-2.51 (4.1)
Equation standard error	0.236	0.230	0.236
Out-of-sample RMSE 2009Q3-2012Q4	0.20	0.22	0.40
Estimation period 1997Q2 – 2008Q2. t-statistics in parentheses.			

Sources: ONS, Markit/CIPS and Bank calculations

## What will go on? Showing Forecast Uncertainty

The last part of the answer to the question “What’s going on?” has to address the question of how good a job the Committee does at anticipating developments in the economy. The MPC was one of the first producers of economic forecasts to try to present an indication of their reliability graphically.<sup>5</sup> The forecast was shown as a sequence of bands round the modal outcome, with the inner band representing the most likely part of the distribution, with a ten per cent probability of realisation, and subsequent bands representing successive deciles. It was assumed that the underlying distribution of forecast errors followed a two-part normal distribution (Britton, Fisher and Whitley, 1998, Wallis, 1999). The bands were constructed by specifying the mode of the distribution and the standard deviations of the normal distributions used for each of the two

<sup>5</sup> Although HM Treasury and the National Institute of Economic and Social Research had, for many years, presented mean absolute forecast errors for their major forecast variables.



components. In 2013, the presentation was changed to show only three bands – with each of these three bands having a thirty per cent chance of realisation. This change, following a suggestion from Ken Wallis which was repeated by David Stockton, was made because the Committee was concerned that the presentation of deciles gave an unrealistic air of precision to the chart. But the Committee has continued to use the two-part normal distribution, the “Zweiseitige Gauss’sche Gesetz” of Fechner (1897), as a framework for specifying the distribution of forecast errors and to produce some graphs showing ten error bands, as I will today.<sup>6</sup>

The assumption that errors are normally distributed is widely made and has many convenient properties. It leads, for example, to the conclusion that the mean is the best summary statistic indicating the location of the distribution, while the standard deviation is the best summary statistic for showing the dispersion of the distribution. As was noted during the financial crisis, it tends to imply that realisations a long way from the mean will occur only rarely. The Committee’s experience was that large forecast errors occurred more frequently than my predecessors had thought likely. In one sense this does not matter; the fan charts represented the subjective judgement of the Committee. One might, nevertheless, hope for some coherence between subjective judgement and subsequent realisation, and the Bank’s Independent Evaluation Office carried out statistical tests which suggested that the forecast errors were distributed materially differently from anything coherent with the *ex ante* judgement of the Committee (Independent Evaluation Office, 2015, p.50).

The t-distribution is a generalisation of the normal distribution, including the latter as a special case (Student, 1908), which requires a third parameter to specify it – the number of degrees of freedom. When the number of degrees of freedom is very large the distribution is very close to the normal distribution; indeed it asymptotically converges to it as the number of degrees of freedom rises to infinity. With a low number of degrees of freedom it has fat tails – representing what seems to be the reality that unusual events are less unusual than they “ought” to be, or rather than a normal distribution would suggest that they are. Fernandez and Steel (1998) show how to construct a two-piece t-distribution using one further parameter which embodies the skewness.<sup>7</sup> So, with the aid of a two-piece t-distribution it is possible to represent both the idea that extreme outcomes are more likely than a normal distribution implies, and the notion that risks may be skewed. The distribution of past forecast errors might provide an indication of the extent both to which the economy is exposed to “unusual” events, and of the extent to which the risks are skewed either up or down.

Chart 4 shows a histogram which displays our actual forecast errors for the level of GDP two years ahead for the forecasts produced from 1998Q1 to 2013Q4. Here, and in all subsequent charts I am showing the errors relative to the modal forecast, because the focus of my discussion is on the distribution around this rather than on the position of the mode. You can see that the errors tended to be clustered between -2 per cent and

---

<sup>6</sup> Wallis (2014) provides a fascinating account of the repeated rediscovery of the two-part normal distribution.

<sup>7</sup> Zhu and Galbraith (2010) describe a two-part t-distribution with different numbers of degrees of freedom in the upper and lower tails. I have not as yet implemented that, but, if the MPC as a whole decides to review the way in which it constructs error bands, I expect it will wish to explore this.

2 per cent. However, the Committee, like everyone else, failed to forecast the recession of 2008/9 and this led to a slew of very large forecast errors. There are in fact six out of sixty-four errors larger than three percentage points and these all happened at the time of the recession.<sup>8</sup> But you can see that, even without these errors, the histogram gives a sense of a downside skew. This largely reflects the disappointing growth the United Kingdom experienced during my first three years on the Committee, as the crisis in the euro area spilled over to the United Kingdom.

There are however, two issues arising in fitting the pattern of forecast errors shown in Chart 4. First of all, the forecast periods overlap with the implication that forecast errors are likely to be serially correlated; I have not made any adjustment for this here. Secondly, the forecast errors for the 2008/9 recession have a very profound influence on the structure of the distribution that can be fitted, notwithstanding that the t-distribution is less affected by outliers than is the normal distribution. It seems to me reasonable to downweight their influence, although I would be reluctant to suppress it completely. What I have done is to construct a synthetic series of forecast errors which comprises three sets of data – the actual series plus two sets of this series with the six largest errors excluded. This means, in principle, that I am assuming the 2008/9 experience comes once in over forty years rather than once in fifteen years. You might think this is excessively optimistic, but the point I want to make is that, even with this treatment, the effects of the experience on the shape of the distribution are pronounced.

This is, however, not the only way of addressing the issue. In particular, the Committee has always said that it has views only on the central ninety per cent of the distribution of errors. That means that the shape of the distribution should not be distorted by observations lying further out. The methods developed by Kapetanios, Mitchell, Price and Fawcett (2015) may make it possible to fit the distribution only to the central ninety per cent, and I expect that the Committee would wish to pursue that should it want to explore further the issues I am discussing here. The results I present are certainly sensitive to the handling of this issue.

In Chart 5 I show a representation of the GDP forecast error probabilities after down weighting the outliers in the way that I have described, and fitting the resulting series to the two-part t distribution.<sup>9</sup> The shaded green areas show ranges, with each shade describing a range of values which has a ten per cent probability. The darkest green shows the ten per cent band closest to the most likely outcome, or mode, and lighter bands show ten per cent bands which are further from the mode. This chart is therefore an analogue of Chart 5.10 published in yesterday's inflation report. But instead of representing the MPC's subjective judgement I am showing a pattern which uses the two-part t distribution to fit our past forecast errors. It is therefore the distribution for our growth forecast at two years which would be appropriate if we thought our future forecasting performance was going to be like our past forecasting performance.


---

<sup>8</sup> A separate concern is that I have treated each of our sixty-four forecasts two years ahead as though they are independent of all other forecasts. That is not the case; the periods overlap. If I were to work with periods which did not overlap I would have only eight observations and it would not be possible to draw any inferences.

<sup>9</sup> I have added to this Chart, and all subsequent fan charts, an offset so that it is centred on zero.


You can see that there is a strong skew to the left. The experience of the recession means that, even after the down weighting I have described, large shortfalls relative to our forecasts are more likely than very high growth figures. It is also the case that the appropriate t-distribution has only 3.6 degrees of freedom. Despite this being a “fat tails” distribution, the consequence is that the central bands are much tighter together than they are if I represent the same error pattern using the MPC’s existing two-part normal distribution. Thus, although the width of the distribution as a whole, which covers ninety per cent of possible outcomes, is marginally wider than that used by the MPC, the central bands are much closer together. Of course, the MPC’s judgement is subjective, so there is no reason why it need match this chart, but I think there is a case for using central bands narrower than those shown in the *Inflation Report*.

**Chart 4: GDP Forecast Errors**



Sources: ONS and Bank calculations

**Chart 5: Fitted GDP Forecast Errors**




Sources: ONS and Bank calculations

Chart 6 and Chart 7 show the analogous results for the inflation forecast errors. You can see that here the skew is to the upside. It is noteworthy that this is despite the experience of the last two years, when the sharp fall in the oil price and reductions in food prices took the inflation rate close to zero. There is, of course, a case for downplaying some of the large upside errors which accrued from 2010-2013, but I think it is less good than for reducing the impact of the recession on the growth distribution. Anyway, I have not made any adjustment here.

The fitted t-distribution has 2.8 degrees of freedom, once again implying a picture rather different from the two-part normal distribution. In this case, the width of the fan is similar for both distributions, but the fitted-distribution has markedly more mass located around the mode. The conclusion that I draw from this exercise is that, although some have been wide of the mark, many of our growth and inflation forecasts are better than we think they are going to be.


Looking at the two charts together, the combination of a downside skew to GDP and an upside skew to inflation might be taken to mean that adverse supply shocks, which depress output and push up on inflation, are more likely than adverse demand shocks, which would pull down on both output and inflation. I am not sure that that is the case. It seems to me that the upward skew of the inflation chart, and in particular the way in which the inflation rate barely dropped below zero during the recession, may instead indicate that there is considerable downward stickiness in prices. That could provide a source of skew in the forecast errors at a time of low inflation.

**Chart 5: Inflation Forecast Errors**



Sources: ONS and Bank calculations

**Chart 4: Fitted Inflation Forecast Errors**

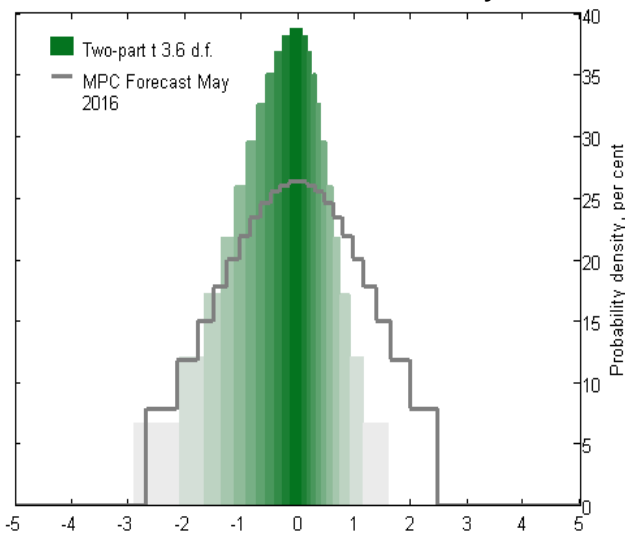


Sources: ONS and Bank calculations

Finally, it is worth making a comparison between what emerges from this and the error distribution assumed by the MPC in its *May Inflation Report*. Chart 8 and Chart 9 repeat the fitted two-part t-distributions with the distributions from the latest *Inflation Report* superimposed. Once again, the modes are located at zero to facilitate the comparison. At least as compared to historic forecasting errors, the MPC's distributions seem to have too much probability mass in the tails. For GDP that seems to be in the upper tail while for inflation it is in the lower tail. This may be desirable because, after all, the *Inflation Report* fan charts are subjective. The comparison between the historic record and what I have just agreed to for the *May Inflation Report* does, nevertheless, raise some questions.

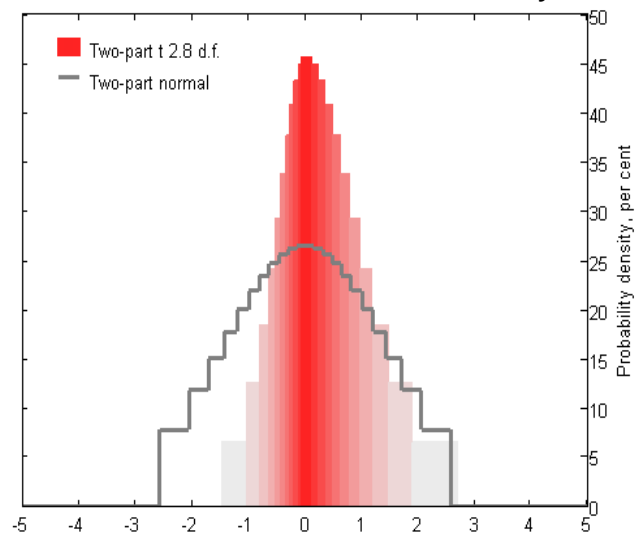
A critique of all of this is, of course, that this historic experience represents a combination of two very different periods, before and after the crisis. One might therefore judge that the experience after the crisis is a better guide to present uncertainty than is the whole of the MPC's record. Equally, it took some time for at least one member of the MPC (me) to realise how things had changed, and that learning process aggravated our forecasting record. Anyway these are issues which others may have to grapple with.

**Chart 6 Fitted GDP Errors and the May IR**



Sources: ONS and Bank calculations

**Chart 7 Fitted Inflation Errors and the May IR**



Sources: ONS and Bank calculations

## Uncertainty is going on

So far I have tried to describe both how we form a picture of what is going on and how we convey the uncertainty surrounding our views of what will go on. But none of you can have avoided noticing that what is going on at the moment is a referendum campaign; we are being asked, for a second time in the voting lives of at least a large minority of us, to vote on the UK's membership of the European Union. Given that the outcome of the vote is going to have important implications for Britain's economic and political future, and since the result is uncertain, you would not be surprised to learn that the Monetary Policy Committee has seen the referendum as a source of uncertainty, with the prospect of this remaining elevated should the nation vote to leave the European Union (Monetary Policy Committee, 2016).

There are any number of reasons why uncertainty should affect economic activity. Keynes (1936) introduced the concept of precautionary saving – people saving up not for some future need, like buying a new car or retirement, but because of what might happen – while Dixit and Pindyck (1994) have drawn attention to the way in which it is likely to affect investment.

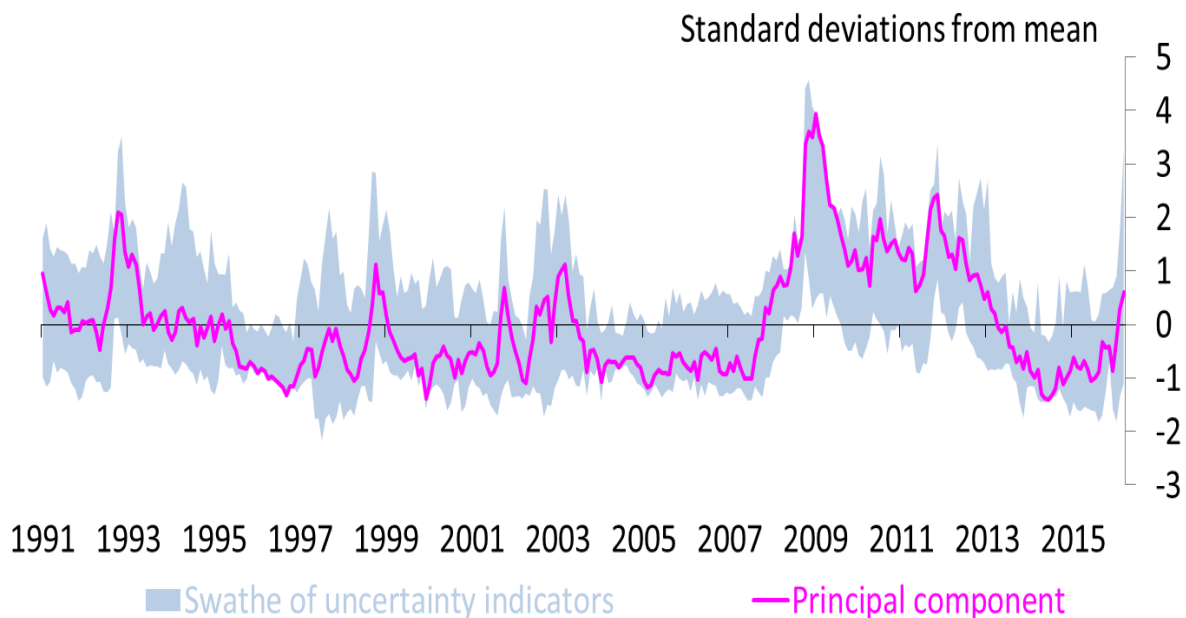
In fact the impact of aggregate uncertainty on household consumption is likely to be rather small. The reason for this is that, for most people, uncertainty about future spending power arises mainly for idiosyncratic reasons, such as changes in family circumstances, rather than changes in the aggregate state of the economy. Since it is the variance, the square of the standard deviation, rather than the standard deviation itself which is likely to affect the growth in consumption (Deaton, 1992), even a doubling of aggregate uncertainty will not have much effect on overall uncertainty. The one change which can have a material impact is an increase in unemployment, but even here, I found a few years ago, using a structural model,

that a two-point increase in unemployment would still have only a relatively small and short-lived impact on saving (Weale, 2012).

In contrast, the impact of uncertainty on investment is likely to be material (Bloom, 2009). Consistent with that, investment is a much more volatile component of demand than is household consumption. Denis and Kannan (2013) explore the effects of uncertainty on the economy of the United Kingdom. They use two measures, the volatility of share prices as implied by the market for options on the FTSE 100 share index and the dispersion of forecasts for GDP growth one year ahead. They suggest that shocks to uncertainty have their maximum effect on GDP at nine months followed by a recovery which takes about two years; not surprisingly, the impact on industrial production is quicker and more pronounced.

Work at the Bank on measuring uncertainty has focused on six indicators, the volatility of the FTSE index, the implied volatility of the exchange rate index deduced from options markets, the dispersion of economic forecasts, a count of media references to uncertainty in an economic context and the balances of household views on unemployment and their financial situation as measured by the GfK monthly survey. None of these is ideal, but the hope is that, by taking their principal component, a reasonable overall picture can be obtained. Chart 8 shows the swathe of these measures together with the principal component. You can see that, although this measure suggests uncertainty has recently picked up to some extent, the level is still appreciably lower than it was between late 2008 and 2012, and not very different from the experience of the Far Eastern crisis of 1998 and the build up to the war in Iraq.

**Chart 8: Swathe of Uncertainty Indicators and their Principal Component**



Sources: Bloomberg, Consensus Economics, Dow Jones Factiva, GfK (research on behalf of the European Commission), Thomson Reuters Datastream and Bank calculations.

I share the collective view of the MPC that the EU referendum is the biggest source of uncertainty at the moment, and this is the main reason why we expect Q2 growth to be weak. There is reasonably clear evidence that the uncertainty has had a material influence on the exchange rate; over the last few months exchange rate movements have tended to occur when news stories about the referendum were dominant. This is the basis for the Committee's assumption that half of the nine per cent fall in the exchange rate since November is associated with uncertainty surrounding the outcome of the referendum. Consistent with a forecast produced on the assumption that government policy is followed, we assume that this component of the fall does not affect the projections for output and inflation.

As in so many areas, the MPC has to produce the best estimates that it can on the basis of the evidence as it is and I am entirely happy with projection that we have just published. Nevertheless, were we to have over-estimated the degree to which the referendum has weighed on sterling, this would raise the prospect of inflation rising noticeably above target in two years or so. This could generate a trade-off between limiting an overshoot to inflation and supporting economic growth. For me that would feel like history repeating itself.

## Conclusions

The last six years have taught me the importance of making the best of what we have. The evidence available to the Monetary Policy Committee is never as good as we would like it to be and we cannot, of course, delay our judgements to wait for statistically solid findings to emerge. Judgements have to be made on the basis of the balance of probabilities rather than firm levels of statistical significance. These points are material for our regular cycle of estimating what is going on ahead of the first official data, interpreting early estimates and explaining what we can and cannot say about the future. What I have shown you here suggests it may be worth taking another look at the way we represent the uncertainty in our forecast.

The referendum raises particular issues of uncertainty. In our forecast we have set out our best estimates of how its effects may be resolved in the event of a vote to remain but it is perfectly possible that events will unfold materially differently from what we forecast even with such a vote. It is important to avoid being seen like the doctor in the *Canterbury Tales* who "knew the cause of everich maladye". Whatever the outcome of the referendum, the Committee will need to set policy to deliver the inflation target, taking due account of the trade-offs between inflation and output variability and with the time taken to return inflation to target reflecting any exceptional circumstances, as the Chancellor's remit letter indicates (Chancellor of the Exchequer, 2016).

## References

- Bell, V, Co, L, Stone, S and Wallis, G (2014).** 'Nowcasting UK GDP Growth'. *Bank of England Quarterly Bulletin 2014 Q1*, pp 58-69.
- Bloom, N (2009).** 'The Impact of Uncertainty Shocks'. *Econometrica*, pp 623-685
- Britton, E, Fisher, P and Whitley, J (1998).** 'The Inflation Report Projections: understanding the Fan Chart'. *Bank of England Quarterly Bulletin 1998 Q1*, pp 30-37.
- Bush, O (2008).** 'Answering Practices Survey of CBI Industrial Trends Survey respondents'. *Confederation of British Industry, April 2008*.
- Chancellor of the Exchequer (2016).** Remit for the Monetary Policy Committee.
- Cunningham, A, Eklund, J, Jeffery, C, Kapetanios, G and Labhard, V (2012).** 'A State Space Approach to Extracting the Signal from Uncertain Data'. *Journal of Business and Economic Statistics*, 30(2), pp 173-180.
- Deaton, A (1992).** 'Understanding Consumption'. *Oxford University Press*.
- Denis, S. and P. Kannan. (2013).** 'The Impact of Uncertainty Shocks on the UK Economy'. IMF Working Paper WP/13/66. <https://www.imf.org/external/pubs/ft/wp/2013/wp1366.pdf>
- Dixit, A and Pindyck, R (1994).** 'Investment under Uncertainty'. *Princeton University Press*.
- Fechner, G T (1897).** 'Kollektivmasslehre Engelmann'. *Leipzig. Edited by G. F. Lipps*.
- Fernandez, C and Steel, M (1998).** 'On Bayesian Modelling of Fat Tails and Skewness'. *Journal of Business and Economic Statistics*, Vol. 8, pp 225–234.
- Independent Evaluation Office (2015).** 'Evaluating Forecast Performance.' *Bank of England*.
- Jones M and McFaddy, M (2003).** 'A Skew Extension of the t- Distribution'. *Journal of the Royal Statistical Society, Series B*, pp 159-174.
- Kapetanios, G, Mitchell, J, Price, S., and Fawcett, N (2015).** 'Generalised forecast density combinations'. *Journal of Econometrics*, Vol 188, pp 150-165.
- Keynes, J M (1936).** 'The General Theory of Employment, Interest, and Money'. *London: Palgrave Macmillan*
- Lui, S, Mitchell, J and Weale, M (2011).** 'Qualitative business surveys: signal or noise?'. *Journal of the Royal Statistical Society, Series A, Part 2*, pp 327-348.
- Mitchell, J (2009).** 'Where are we now? The UK recession and nowcasting GDP growth using statistical models'. *National Institute Economic Review*, No. 209, pp 60-69.
- Monetary Policy Committee (2016).** 'Inflation Report'. May 2016.
- Stratford, K (2013).** 'Nowcasting World GDP and Trade using Global Indicators'. *Bank of England Quarterly Bulletin 2013 Q3*, pp 233-243.
- Student (1908).** 'The Probable Error of a Mean'. *Biometrika* 6 (1), pp 1-25



**Wallis, K F (1999).** 'Asymmetric Density Forecast of Inflation and the Bank of England Fan Chart'. *National Institute Economic Review*, January, pp 106-112.

**Wallis, K F (2014).** 'The Two-Piece Normal, Binormal, or Double Gaussian Distribution: Its Origin and Rediscoveries'. *Statistical Science*, Vol 29, pp 106-112.

**Walton, A (2016).** 'International Comparison of GDP Revisions'. *Office for National Statistics*.

**Weale, M R (2012).** 'From retailers' paradise to shoppers' strike: what lies behind the weakness in consumption?'. Speech at Cass Business School, 29<sup>th</sup> February.

**Wheeler, T (2010).** 'What can we learn from surveys of business expectations?'. *Bank of England Quarterly Bulletin* 2010 Q3, pp 190-98.

**Zhu, D and Galbraith, J (2010).** 'A Generalized Asymmetric t-distribution with Application to Financial Econometrics'. *Journal of Econometrics*, Vol 157, pp 297-305.

## Appendices

### A. A simple model of measurement error in preliminary GDP

The regression coefficient below 1 on the initial estimate in an equation explaining the final estimate can be understood in the following way. Suppose that the economy consists of two sectors, a and b. Sector a has a weight of  $w$  and sector b a weight of  $1-w$ . The output growth of sector a measured relative to mean output growth is, with time subscripts suppressed,

$$y_a = w\varepsilon + \varepsilon_a,$$

while the output of sector b is

$$y_b = (1 - w)\varepsilon + \varepsilon_b.$$

Thus the output of each sector consists of a common component and an idiosyncratic component. These cannot be distinguished in the data for any individual quarter, but the idiosyncratic components of the two sectors are uncorrelated.

Total output is

$$y = y_a + y_b = \varepsilon + \varepsilon_a + \varepsilon_b.$$

However, only sector a is observed initially, with the data for sector b coming in later. The initial estimate is therefore

$$y^* = \varepsilon + \varepsilon_a/w$$

If the final estimate is regressed on the initial estimate, the regression coefficient is  $\beta = w \frac{w + \frac{\sigma_a^2}{\sigma^2}}{w^2 + \frac{\sigma_a^2}{\sigma^2}}$ , which tends

to 1 if there is no idiosyncratic component in sector a. It does not matter if there is still an idiosyncratic component in sector b.

Suppose that the initial estimates come from 40% of the economy, so that  $w=0.4$ . Then if the ratio  $\frac{\sigma_a^2}{\sigma^2} = 1/3$  the resulting regression coefficient will be  $0.4 \times 0.73 / 0.39 = 0.75$ . Thus a substantial but not overwhelming value for  $\sigma_a^2$  will be enough to deliver the sort of coefficient observed in practice.

## B. The two-piece t-distribution

The two-piece t- distribution has the following density function. Here  $G()$  indicates a gamma-function,  $\mu$  is the mode of the two-part distribution and, when there is no skew, it is also the mean and the median,  $\sigma$  is a measure of dispersion which, in the absence of skew is also the standard deviation,  $\theta$  is the number of degrees of freedom and  $\gamma$  determines the skew.  $H()$  is the Heaviside function which takes a value 1 when its argument is zero or positive and 0 when its argument is negative.

$$p(x_i|\mu, \sigma, \theta, \gamma) = \left( \frac{2}{\gamma + \frac{1}{\gamma}} \right) \frac{G\left(\frac{\theta+1}{2}\right)}{G\left(\frac{\theta}{2}\right) (\pi\theta)^{1/2}} \sigma^{-1} \left[ 1 + \frac{(x_i - \mu)^2}{\theta\sigma^2} \left\{ \frac{H(x_i - \mu)}{\gamma^2} + \gamma^2 H(\mu - x_i) \right\} \right]^{-(\theta+1)/2}$$

The density function becomes symmetric when  $\gamma$  is set to 1. A value greater than this generates a skew to the right while a value below this generates a skew to the left. The t- distribution converges to the normal distribution as  $\theta$  tends to infinity. In practice a normal distribution can be generated by setting  $\theta$  to some reasonably large number like two hundred. Thus restricting  $\gamma$  and restricting  $\theta$  allow us to explore the special cases of symmetry and normality (with or without symmetry).

The density function is fitted to a sequence of observations  $\{x_j\}$  by maximum likelihood. That is, values of the parameters,  $\mu, \sigma, \theta$  and  $\gamma$  are chosen to maximise the sum of the logarithms of the probabilities of each observation conditional on the parameter set.

$$\text{Max}_{\mu, \sigma, \theta, \gamma} \sum_i \log(p(x_i|\mu, \sigma, \theta, \gamma))$$

This maximisation problem can be solved in MATLAB using the routine *fminunc*.